Universal Clustering via Crowdsourcing

Ravi Kiran Raman and Lav R. Varshney

arXiv:1610.02276v1 [cs.HC] 5 Oct 2016

Abstract

Consider unsupervised clustering of objects drawn from a discrete set, through the use of human intelligence available in crowdsourcing platforms. This paper defines and studies the problem of *universal clustering* using responses of crowd workers, without knowledge of worker reliability or task difficulty. We model stochastic worker response distributions by incorporating traits of memory for similar objects and traits of distance among differing objects. We are particularly interested in two limiting worker types—temporary workers who retain no memory of responses and long-term workers with memory. We first define clustering algorithms for these limiting cases and then integrate them into an algorithm for the unified worker model. We prove asymptotic consistency of the algorithms and establish sufficient conditions on the sample complexity of the algorithm. Converse arguments establish necessary conditions on sample complexity, proving that the defined algorithms are asymptotically order-optimal in cost.

Index Terms

Budget optimality, clustering, crowdsourcing, universal information theory, unsupervised learning

I. Introduction

Crowdsourcing has grown in recent times as a potent tool for performing complex tasks using human skill and knowledge. It is increasingly being used to collect training data for novel machine learning problems. Almost *a fortiori*, there is no prior knowledge on the nature of the task and so the use of general human intelligence has been needed [1]. As such, this setting requires processing human-generated signals in the absence of prior knowledge about their properties; this setting requires universality.

Because crowdsourcing often employs unreliable workers, the signals they generate are inherently noisy [2]. Hence, responses of crowd workers are modeled as outputs of a noisy channel. Although these channels are unknown to the employer, crowdsourcing techniques have thus far made assumptions on either the channel distribution or structure to design appropriate decoders. We define an alternative approach—*universal crowdsourcing*—that designs decoders without channel knowledge, and develop achievability and converse arguments that demonstrate order-optimality.

The emergence of diverse online crowdsourcing platforms such as Amazon Mechanical Turk and Upwork (formerly oDesk) has created the option of choosing between temporary workers and long-term workers. That is, tasks can be completed by either soliciting responses from a large number of workers performing small parts of a large task, or a specialized group of employees who work long-term on the task at hand.

The trade-off between reliability and cost of each type of worker pool warrants systematic study. Whereas temporary workers are inexpensive and easily available, some labor economists argue the excess cost of long-term employment is worthwhile due to the reliability and quality of work it ensures [3]. However, no quantitative characterization for this conjecture exists. As we will see, the results of this paper allow such comparisons.

Since workers are human, they are subject to several factors that affect human decision making, as identified in behavioral economics [4].

A standard assumption in crowdsourcing and in universal clustering [5], [6] has been independent and identically distributed (i.i.d.) worker responses across time/tasks. However, empirical evidence argues against temporal independence for individual worker responses [7]. Due to the *availability heuristic*, worker responses may rely on the immediate examples that come to a person's mind, indicating memory in responses across tasks/time. Further, this influence may be more due to salient (vivid) information rather than full statistical history.

Due to the *anchoring and adjustment heuristic*, people tend to excessively rely on a specific trait of an object in decision making and further due to the *representativeness heuristic*, people tend to assume commonality among

This work was supported in part by NSF Grants IIS-1550145 and CCF-1623821, and by Air Force STTR Contract FA8650-16-M-1819.
The authors are with the University of Illinois at Urbana-Champaign.

objects. These traits indicate there is a notion of distance among the response distributions corresponding to different objects.

To capture memory and distance, we define a unified model of worker responses, and then study two limiting cases. First we consider responses of temporary workers who respond independently across tasks and across workers. We then consider responses of long-term workers with object-specific memory. Specifically, we consider a Markov memory model wherein the response to an object is dependent on the most recent response and the response to the most recent occurrence of an object of the same class; generalizations to other Markov models follows readily. In both cases, we address questions of universality and sample complexity for reliable clustering, providing benchmarks for worst-case performance.

A. Prior Work

There is a vast and rich literature on crowdsourcing and clustering; we describe a non-exhaustive listing of particularly relevant prior work.

Algorithm design for crowdsourcing typically focuses on minimizing the cost of reliability. In particular, algorithms with order-optimal budget-reliability trade-offs have been designed for binary classification with unknown (but i.i.d.) crowd reliabilities [5]. Efficient algorithms for multi-class labeling have been proposed in [8] albeit without cost optimality guarantees. More recently, non-parametric permutation models of crowd workers were considered for binary classification [9].

An alternate strategy for multi-class labeling with workers lacking sufficient domain expertise is to decompose the overall task into simpler subtasks, and introduce redundancy through an error control code [10]. This approach is implicitly effective for mismatched crowdsourcing for speech transcription [11].

Separate from crowdsourcing, the problem of clustering has been widely studied. Algorithms such as k-means clustering and its generalization to other Bregman divergence similarity measures [12] are popular methods that incorporate distance-based clustering.

The problem of universal clustering was considered in a communication setting [6], [13], such that messages communicated across an unknown channel, after encoding using a random codebook, are clustered by exploiting dependency among outputs of similar messages. Particularly the decoder uses the minimum partition information functional [14] to perform optimal clustering. Similar information-based agglomerative clustering schemes have also been explored [15].

Classification using crowdsourced responses in a clustering framework, followed by a labeling phase performed by a domain expert has been studied experimentally [16].

B. Main Contributions

In this work, we provide a theoretical study of universal crowdsourcing for clustering. In particular, we focus on the design of universal clustering algorithms with provable asymptotic consistency and order optimality in sample complexity.

The presence of memory in worker response demands an approach that differs from past crowd algorithms defined for i.i.d. models [5]. Notwithstanding [10], in the crowdsourcing framework herein, we do not have the opportunity to encode messages. This calls for a treatment different from past work in universal clustering [6].

We first consider the case of temporary workers without memory, wherein, we design a distance-based universal clustering decoder that uses distributional identicality among objects of the same class. That is, two objects are clustered together if the f-divergence between the conditional distributions of the responses is small. We prove asymptotic consistency of the algorithm and prove order optimality in sample complexity. The algorithm applies directly to a large class of similarity measures.

We then consider long-term employees with object-specific memory. Specifically we consider a Markov memory model wherein the response to an object is dependent on the most recent response and the response to the most recent occurrence of an object of the same class. For this model, we show the existence of information-based universal clustering strategies that perform asymptotically reliable clustering. Further, we study the sample complexity of the decoder and show order optimality by comparing with the necessary cost. We also highlight the extension of the algorithm to higher-order Markov memory models. We also observe that the universal clustering algorithm is structurally similar to traditional clustering algorithms such as the MIRN (mutual information relevance networks)

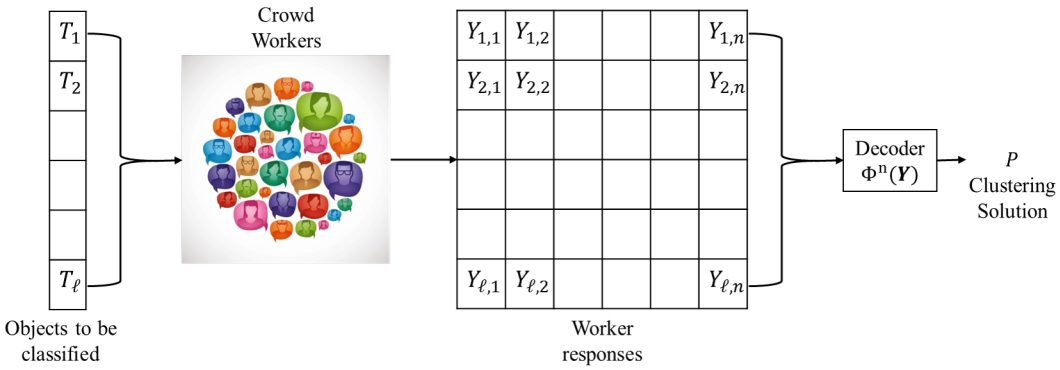

Fig. 1. Model of the crowdsourcing system

clustering [17], and the minimum partition information clustering [14] algorithms under added constraints on the channel model.

Finally, we use results obtained for these two limiting cases to construct a clustering algorithm for a unified worker model. We prove asymptotic consistency. Further, we show order optimality of the sample complexity as a function of the number of objects to be clustered.

II. Model

A. System

We formalize the model of the crowdsourcing system by first describing a unified worker model and then specializing to the cases of workers with and without memory. We aim to design universal decoders that cluster a given set of objects using the crowd responses. For any index vector S, let Z_s be the set $\{Z_i : i \in S\}$, where Z is any vector. Similar notation shall be used for a matrix of indices S and matrix $\mathbf{Z}$.

We consider the problem of crowd workers employed to perform classification of objects. For instance, consider the task of classifying images of dogs according to their breeds. The workers observe images and respond with the breed of the dog in the image. Since worker responses are noisy, in the absence of knowledge of worker channels it is not feasible to identify the labels (breeds of dogs) accurately. Thus, we aim to cluster the dogs according to their breeds and determine the labels of each cluster by using a domain expert. The crowdsourcing system model is depicted in Fig. 1.

Let $\mathcal{T}$ be a finite alphabet of object clusters. Without loss of generality, let us assume that $\mathcal{T} = [\tau] = \{1, \ldots, \tau\}$, where $\tau < \infty$ is a constant. Each object is viewed by crowd workers as $X \in \mathcal{X}$, which has some relation to its type $T \in \mathcal{T}$.

Let the set of objects to be clustered be $\{X_1, \ldots, X_\ell\}$. Let the label of object X_i be T_i, for all $i \in [\ell]$. That is, the objects to be clustered are treated as manifestations of the various object labels. Thus clustering the set of objects is the same as clustering according to their object labels. Let us assume that the objects are drawn according to an unknown prior $P_T(\cdot)$ on the set of classes.

For each object $X_i, i \in [\ell]$, the crowdsourcing system solicits n responses, Y_i^n, from workers employed to classify the objects according to their labels. The collection of responses is given by the matrix $\mathbf{Y} = (Y_1^n, \ldots, Y_\ell^n)' = [Y_{i,j}]_{i\in[\ell],j\in[n]}$. Let $W \in \mathcal{W}$ be the index of a worker and let $S_W \subset [\ell] \times [n]$ be the index set corresponding to the responses offered by worker W in $\mathbf{Y}$. We assume that $S_w \cap S_{\hat{w}} = \emptyset$ for any two workers $w \neq \hat{w}$ and $\cup_{w\in\mathcal{W}} S_w = [\ell] \times [n]$.

Let $\mathcal{Q}$ be the set of all conditional probability mass functions (pmfs) that characterize worker responses. Then, $Y_{S_W} \sim Q^{(W)}(Y_{S_W}|T_{\hat{S}_W})$ where $Q^{(w)}(\cdot) \in \mathcal{Q}$ is the distribution characterizing the response of worker $w \in \mathcal{W}$ and $\hat{S}_W$ is the corresponding set of indices of objects. Models of these distributions are detailed later.

With regard to the example of classifying dog images, the object label, T, is the breed of the dog in an image; the object, X, is the image of a particular dog; and the responses of the crowd workers Y^n are the breeds they categorize the image to.

Since clustering is performed solely based on the response of the workers, it is fair to assume that the number of clusters that can be formed is directly dependent on $|\mathcal{Y}|$. However, it is not essential for every worker to answer every question in a practical crowdsourcing platform. Thus we assume that the workers either respond with an answer in $\mathcal{T}$ or offer a 'null' response ξ to every task [18]. Thus, without loss of generality, we assume that $Y_{ij} \in \mathcal{Y} = \mathcal{T} \cup \{\xi\}$, for all $i \in [\ell], j \in [n]$.

B. Universal Clustering Performance

Definition 1 (Correct Clustering): A *clustering* of a set of objects $X_1, \ldots, X_\ell$ is a partition P of $[\ell]$. The sets of a partition are referred to as *clusters*. The clustering is said to be *correct* if

$$T_i, T_j \in C \iff T_i = T_j,$$

for all $i, j \in [\ell]$, $C \in P$. For a given set of object labels, T^ℓ, let $P^*\left(T^\ell\right)$ be the correct clustering.
Let $\mathcal{P}$ be the set of all partitions of $[\ell]$.

Definition 2 (Partition Ordering): A partition P is *finer* than P', if the following ordering holds

$$P \preceq P' \iff \text{ for all } C \in P, \text{ there exists } C' \in P' : C \subseteq C'.$$

Similarly, a partition P is said to be *denser* than P' if $P \succeq P' \iff P' \preceq P$.

Definition 3 (Universal Clustering Decoder): A *universal clustering decoder* is a sequence of functions $\Phi^{(n)} : \mathbf{Y} \to \mathcal{P}$ that are designed in the absence of knowledge of $\mathcal{Q}$ and P_T. Here the index n corresponds to the number of crowd responses collected per object.

We now define how to characterize decoder performance.

Definition 4 (Error Probability): Let $\Phi^{(n)}(\cdot)$ be a universal decoder. Then, the *error probability* is given by

$$\begin{aligned} P_e(\Phi^{(n)}) &= \Pr\left[\Phi^{(n)}(\mathbf{Y}) \neq P^*\left(T^\ell\right)\right] \\ &= \mathbb{E}_{P_T^{\otimes \ell}}\left[\mathbb{E}\left[\mathbb{1}\left\{\Phi^{(n)}(\mathbf{Y}) \neq T^\ell\right\} | T^\ell\right]\right], \end{aligned} \tag{1}$$

where $P_T^{\otimes \ell}(t^\ell) = \prod_{i=1}^{\ell} P_T(t_i)$ and $\mathbb{1}\{\cdot\}$ is the indicator function.

Definition 5 (Asymptotic Consistency): A sequence of decoders $\Phi^{(n)}$ is said to be *universally asymptotically consistent* if

$$\lim_{n \to \infty} P_e(\Phi^{(n)}) = 0, \text{ for all } P_T \in \mathcal{M}(\mathcal{T}),$$

where $\mathcal{M}(\cdot)$ is the space of all prior distributions on the set of objects, $\mathcal{T}$.

Definition 6 (Sample Complexity): Let $\epsilon > 0$ be the permissible error margin. Then the *sample complexity* of the universal clustering problem is

$$N^*(\epsilon) = \min\left\{n \in \mathbb{N} : \max_{P_T \in \mathcal{M}(\mathcal{T})} P_e(\Phi^{(n)}) < \epsilon\right\},$$

where the minimum is taken over the set of all sequences of universal decoders $\Phi^{(n)}$.

For simplicity, we will use Φ to denote $\Phi^{(n)}$ when it is clear from context.

C. Workers

We now define a model to characterize crowd worker responses. Let us assume the crowdsourcing system employs n crowd workers chosen at random from $\mathcal{W}$. Without loss of generality, let us assume that the set of workers chosen is $[n]$. We assume that each worker responds to every object $X_i, i \in [\ell]$ and the set of responses of worker j is $\{Y_{1,j}, \ldots, Y_{\ell,j}\}$. We assume that the responses of each worker $w \in \mathcal{W}$ are drawn according to the conditional distributions Q^w.

We assume that the responses of each worker depend on prior responses in a Markov sense. In particular, define the set of neighbors as $\mathcal{N}_i = \{\{i-1\} \cup \{\max\{k \in [i-1] : T_k = T_i\}\}\}$ for any $i \in [\ell]$ i.e., the most recent object

and the most recent occurrence of an object of the same class. The response of any worker $j \in [n]$, for any $k \leq \ell$ satisfies

$$\Pr\left[(Y_{1,j}, \ldots, Y_{k,j}) = (y_1, \ldots, y_k) | T^k = t^k\right] = \prod_{i=1}^{k} Q^{(j)}(y_i | y_{\mathcal{N}_i}, t_i). \tag{2}$$

Additionally we assume that for any worker $j \in [n]$ and $i \leq \ell$, for every $t \in \mathcal{T}$,

$$\Pr\left[Y_i = y | T_i = t\right] = Q^{(j)}\left(y|t\right). \tag{3}$$

That is, the worker responses are dependent on the prior responses such that the marginal conditional distribution of the response, given an object is identical across objects of the same class (that is, the marginals are invariant across permutations of the given set of objects).

We also assume that, given the object, the responses are independent across workers. That is,

$$\Pr\left[(Y_{i,1}, \ldots, Y_{i,k}) = (y_1, \ldots, y_k) | T_i = t\right] = \prod_{j=1}^{k} \Pr\left[Y_{i,j} = y_j | T_i = t\right]. \tag{4}$$

Thus, the unified worker model of crowd responses is characterized by (2), (3), and (4).

As mentioned in Section I, in addition to the unified worker model, we focus on two special classes of workers—temporary workers without memory and long-term workers with memory.

In particular, when temporary workers are employed we assume they do not retain memory of their prior responses and so are independent across time. In order to solicit such responses, we may assume that the crowdsourcing platform delegates each task to a sequence of n workers selected uniformly at random from a sufficiently large crowd to ensure independence of responses across objects. That is, if the index of responses of worker w is $S_w \subset [\ell] \times [n]$, then

$$\Pr\left[\mathbf{Y}_{S_w} = y_{S_w} | T^\ell = t^\ell\right] = \prod_{(i,j) \in S_w} Q^{(w)}(y_{i,j} | t_i). \tag{5}$$

Further, the responses also satisfy (4).

The second special class of workers we consider is long-term workers with Markov memory. In particular, responses of this class of workers are characterized by (2) and (4).

III. Temporary Workers

Consider the scenario of clustering using responses of temporary workers (without memory). As mentioned in the system model, we assume worker responses are independent across objects and workers.

We first introduce f-divergences and their properties, so as to characterize the quality of crowd responses. We then define the universal clustering algorithm and prove asymptotic consistency and order optimality in sample complexity.

A. f-Divergence

To measure the separation among the conditional distributions of crowd responses to different object classes, we use the Csiszár f-divergence [19], [20].

Definition 7 (f-divergence): Let p, q be discrete probability distributions defined on a space of m alphabets. Given a convex function $f : [0, \infty) \to \mathbb{R}$, the *$f$-divergence* is defined as:

$$D_f(p\|q) \triangleq \sum_{i=1}^{m} q_i f\left(\frac{p_i}{q_i}\right). \tag{6}$$

The function f is said to be *normalized* if $f(1) = 0$.

Some specific f-divergences are the KL divergence $D(p\|q)$ and the total variational distance $\delta(p, q)$. Specifically, the KL divergence and the total variational distance are the f-divergences corresponding to the functions $f(x) = x \log x$ and $f(x) = |x - 1|$ respectively. We now state some bounds for f-divergences.

Theorem 1 ([21, Chapter II.1]): Let p, q be discrete probability distributions on a space of m alphabets such that there exist r, R satisfying $0 \leq r \leq \frac{p_i}{q_i} \leq R \leq \infty$ for all $i \in [n]$. Let $f : [0, \infty) \to \mathbb{R}$ be a convex and normalized function satisfying the following criteria:

1) f is twice differentiable on $[r, R]$, and
2) there exist real constants $c, C < \infty$ such that

$$c \leq xf''(x) \leq C, \text{ for all } x \in (r, R).$$

Then, we have,

$$cD(p\|q) \leq D_f(p\|q) \leq CD(p\|q). \tag{7}$$

For ease, we refer to the constraints on f in Theorem 1 as *smoothness* constraints. If f is twice differentiable in $[r, R]$, then we know that there exists a constant L such that f is L-Lipschitz.

Theorem 2 ([21, Chapter II.3]): Let $f : [0, \infty) \to \mathbb{R}$ be convex, normalized, and L-Lipschitz on $[r, R]$. Then,

$$0 \leq D_f(p\|q) \leq L\delta(p, q). \tag{8}$$

Further, Pinsker's inequality lower bounds the KL divergence $D(p\|q)$ with respect to the total variational distance as

$$D(p\|q) \geq (2 \log_2 e)\delta^2(p, q). \tag{9}$$

Corollary 1: For any convex and normalized function f that satisfies the smoothness constraints and is L-Lipschitz,

$$\kappa\delta^2(p, q) \leq D_f(p\|q) \leq L\delta(p, q), \tag{10}$$

where $\kappa = 2c \log_2 e$.

Proof: The result follows from Theorems 1 and 2, and (9). ∎

These inequalities are used to prove consistency of the designed universal clustering decoder.

B. Task Difficulty for Worker Pool

Let $Q_1, \ldots, Q_\tau$ be the conditional response distributions given the object class, defined as

$$Q_i(y) \triangleq \Pr[Y = y|T = i] = \mathbb{E}\left[Q^{(W)}(Y = y|T = i)\right],$$

where the expectation is taken over the workers in the pool. Since the responses are obtained from temporary workers chosen at random, it suffices to consider these expected conditional response distributions.

Definition 8 (Distance Quality): For a given pool of temporary workers, the *difficulty of the tasks* is defined as

$$\theta_d \triangleq \min_{\{i,j \in \mathcal{T}, i \neq j\}} D_f(Q_i\|Q_j), \tag{11}$$

where D_f corresponds to the f-divergence chosen as the notion of similarity for the problem at hand.
The operational significance of this informational definition of distance (Definition 8) will emerge in coding theorems Lemma 1 and Theorem 3.

Clustering is performed using the maximum likelihood estimates of the f-divergence between distributions corresponding to responses to objects. Convergence of the empirical estimates is asymptotically consistent and the rates of convergence are discussed in Appendix A.

C. Universal Clustering using Temporary Workers

Responses to objects of the same class are identical in distribution. Thus, we perform universal clustering, $\Phi_{\text{temp}}(\mathbf{Y})$, according to Algorithm 1. That is, the algorithm identifies the cliques in the graph obtained by thresholding the f-divergence between the corresponding empirical distributions. The functioning of the algorithm is depicted in Fig. 2.

Lemma 1: For $f(x) = |x - 1|$ and $\gamma_n = c_1 n^{-\alpha}$, $\alpha \in [0, 1/2]$ let $\gamma_n < \theta_d/2$. For any other convex function f satisfying the smoothness constraints let

$$\gamma_n = c_1 n^{-\beta} < \frac{\kappa}{2L}\theta_d^2 + 2\frac{\kappa}{L^2}\left(1 - \sqrt{1 + \frac{L\theta_d^2}{2}}\right),$$

Algorithm 1 Clustering with temporary workers, $\Phi_{\text{temp}}(\mathbf{Y})$

if $f(x) = |x-1|$ **then**
  $\gamma_n \leftarrow c_1 n^{-\alpha}$, where c_1 is a constant and $\alpha \in [0, 1/2]$
else
  $\gamma_n \leftarrow c_1 n^{-\beta}$, where c_1 is a constant and $\beta \in [0, 1]$
end if
Determine empirical distributions q_j, $j \in [\ell]$
Construct $G = ([\ell], E)$, s.t. $(i,j) \in E \iff D_f(q_i, q_j) \leq \gamma_n$
$\mathcal{C} = \{C : C \text{ is a maximal clique in } G\}$
Select minimal weight partition of $[\ell]$ from $\mathcal{C}$

where $\kappa = 2c \log_2 e$ and the function f is L-Lipschitz.

Define the ball of radius ρ centered at p as

$$B_f(p, \rho) = \{q : D_f(p\|q) \leq \rho\}.$$

If for all $i \in [\ell]$, the empirical distribution of responses $q_i \in B_f(Q_{t_i}, \gamma_n/2)$ then, $\Phi_{\text{temp}}(\mathbf{Y}) = P^*(t^\ell)$, the correct clustering of the set of objects.

Proof: Let us first consider $f(x) = |x-1|$. Since $q_i \in B_f(Q_{t_i}, \gamma_n/2)$ for all $i \in [\ell]$, we have

$$\max_{\{i,j \in [\tau]: T_i = T_j\}} \delta(q_i, q_j) \leq \gamma_n \leq \frac{\theta_d}{2} \leq \min_{\{i,j \in [\tau]: T_i \neq T_j\}} \delta(q_i, q_j).$$

Let $C_i = \{j \in [\ell] : T_j = i\}$, for any $i \in [\tau]$. Then, for $j, k \in C_i$, $D_f(q_j, q_k) < \gamma_n$ and so $(j,k) \in E$. Thus, C_i is a clique of G.

Further, this observation also implies that for any $i \in C_t, j \in C_{t'}, t \neq t'$, $D_f(q_i, q_j) > \gamma_n$ and so $(i,j) \notin E$. Thus, any set $C \subseteq [\ell]$ such that there exist $i, j \in C$, with $T_i \neq T_j$, is not a clique in G.

Thus C_i is a maximal clique in G for all $i \in [\tau]$. Thus $\Phi_{\text{temp}}(\mathbf{Y}) = \{C_i : i \in [\tau]\} = P^*(t^\ell)$, the correct partition.

For the second part of the lemma, from (10), we note that the condition on γ_n guarantees

$$\max_{\{i,j \in [\tau]: T_i = T_j\}} D_f(q_i\|q_j) \leq \gamma_n < \min_{\{i,j \in [\tau]: T_i \neq T_j\}} D_f(q_i\|q_j).$$

Thus the result follows from a very similar argument. ■

From Lemma 1, we observe that when the empirical distributions are sufficiently close to the corresponding true distributions, the algorithm outputs the correct clustering. Using this result, we prove consistency of the algorithm.

Theorem 3: If $f(x) = |x-1|$, then for any $\alpha \in (0, 1/2)$ and constant c_1, for

$$n \gtrsim \max\left\{ \left(\frac{2c_1}{\theta_d}\right)^{1/\alpha}, \left(\frac{4 \log \ell}{c_0 c_1^2}\right)^{1/(1-2\alpha)} \right\}, \tag{12}$$

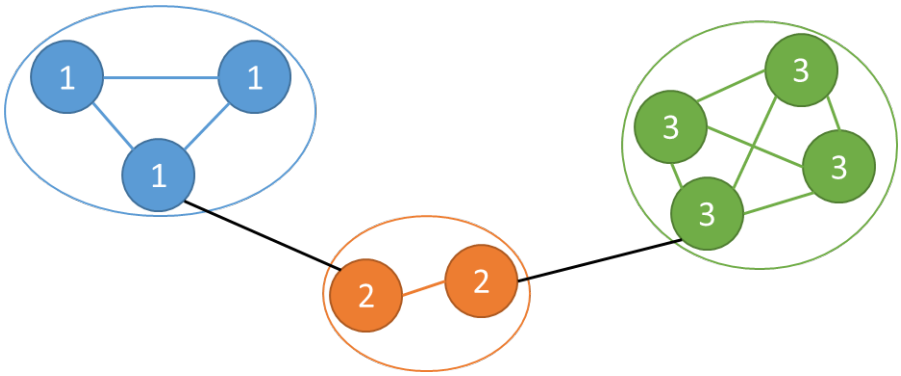

Fig. 2. Distance-based clustering of 9 objects of 3. The graph is obtained by thresholding the f-divergences of empirical distributions of responses to objects. The clustering is then done by identifying the maximal clusters in the thresholded graph.

sufficiently large, $\Phi_{\text{temp}}(\cdot)$ achieves arbitrarily low clustering error probability. For fixed ℓ and θ_d, it is universally asymptotically consistent.

For any other convex, normalized function f satisfying the smoothness constraints, for any $\beta \in (0,1)$ and constant c_1, for

$$n \gtrsim \max\left\{\left(\frac{2c_1L^2}{\kappa}\mu_{\theta_d}\right)^{1/\beta}, \left(\frac{C\log l}{c_1}\right)^{1/(1-\beta)}\right\}, \tag{13}$$

where,

$$\mu_{\theta_d} = \left(L\theta_d^2 + 4\kappa\left(1-\sqrt{1+\frac{L\theta_d^2}{2}}\right)^{-1}\right),$$

with sufficiently large constant c_1, $\Phi_{\text{temp}}(\cdot)$ achieves arbitrarily low clustering error probability. For fixed ℓ and θ_d, it is universally asymptotically consistent.

Proof: For $f(x) = |x-1|$ and $n \geq \left(\frac{2c_1}{\theta_d}\right)^{1/\alpha}$, $\gamma_n \leq \theta_d/4$. Thus, when $f(x) = |x-1|$, we can bound the error probability as follows

$$\begin{aligned} P_e(\Phi_{\text{temp}}) &\leq \mathbb{E}_{P_T^{\otimes l}}\left[\Pr\left[\exists i \in [l] : q_i \notin B(Q_{t_i}, \gamma_n/2)\right]\right] \\ &\leq \ell\left[(n+1)^{|\mathcal{Y}|}\exp\left(-c_0 n\frac{\gamma_n^2}{4}\right)\right] \\ &= \exp\left(\left(\log\ell + (\tau+1)\log(n+1) - \frac{c_0c_1^2}{4}n^{1-2\alpha}\right)\right), \end{aligned} \tag{14}$$

where (14) follows from the union bound and Lemma 5. Thus, the cost conditions given in (12) and asymptotic consistency follow.

Using a similar argument and Lemma 1, we obtain (13). ■

Corollary 2: Given $\mathcal{T} = [\tau]$ with $\tau < \infty$ a constant:

1) for a constant $\theta_d > 0$, $N^*_{\text{temp}}(\epsilon) = O\left((\log\ell)^{1/(1-\beta)}\right)$ and taking $\beta \to 0$, we observe that $N^*_{\text{temp}}(\epsilon) = O(\log\ell)$ for any of the similarity metrics;
2) for a constant ℓ, for $f(x) = |x-1|$, $N^*_{\text{temp}}(\epsilon) = O\left(\theta_d^{-1/\alpha}\right)$ and taking $\alpha \to 1/2$, $N^*_{\text{temp}}(\epsilon) = O(\theta_d^{-2})$. On the other hand, for other convex functions f satisfying the smoothness constraints, $N^*_{\text{temp}}(\epsilon) = O(\theta_d^{-1/\beta})$ and specifically, taking $\beta \to 1$, $N^*_{\text{temp}}(\epsilon) = O(\theta_d^{-1})$.

Proof: The results follow directly from Theorem 3. ■

To summarize, this subsection has defined a universal clustering algorithm that is asymptotically consistent and also described sufficient conditions on sample complexity of universal clustering.

D. Lower Bound on Sample Complexity

We now show matching lower bounds on the sample complexity for universal clustering.

Theorem 4: Let $\min_{i,j\in[\tau]}\delta(Q_i,Q_j) = \theta_d$. Then the sample complexity of universal clustering satisfies

$$N^*_{\text{temp}} = \Omega\left(\frac{\log\ell - \log\epsilon}{\theta_d^2}\right).$$

Proof: Consider the prior distribution P_T such that $P_T(1) = P_T(2) = 0.5$. Let ψ_{ij} be the binary hypothesis testing problem given by:

$$\psi_{ij} : \begin{cases} H_0 : T_i = T_j \\ H_1 : T_i \neq T_j \end{cases}. \tag{15}$$

There exists $\binom{\ell}{2}$ such binary hypothesis tests. Choose a set of tests $\tilde{\psi} \subset \{\psi_{ij} : i,j \in [\ell], i \neq j\}$ of cardinality $\ell/2$ such that no two tests in the set share a common object. This indicates that the binary hypothesis tests are independent of each other owing to the independence across objects.

Let Φ be a decoder for the clustering problem. Then a correct solution to the clustering problem implies a correct solution to ψ_{ij}, for all $i, j \in [\ell], i \neq j$. This implies that an instance of correct clustering translates to correct decisions in all tests in $\tilde{\psi}$. Thus,

$$P_e(\Phi) \geq 1 - \prod_{\{i,j\in[\ell]:\psi_{ij}\in\tilde{\psi}\}} (1 - \Pr[\text{error in } \psi_{ij}]) \tag{16}$$

$$\geq 1 - \prod_{\{i,j\in[\ell]:\psi_{ij}\in\tilde{\psi}\}} \left(1 - \frac{1}{4}\exp(-2nB_{ij})\right) \tag{17}$$

$$\geq 1 - \left(1 - \frac{1}{4}\exp(-2nB_{\max})\right)^{\lfloor \ell/2 \rfloor} \tag{18}$$

$$= \frac{\lfloor \ell/2 \rfloor}{4}\exp(-2nB_{\max}) + o\left(\exp(-4nB_{\max})\right), \tag{19}$$

where (16) follows from the independence of the binary tests and (17) follows from the Kailath lower bound [22]. Here B_{ij} is the Bhattacharyya distance corresponding to the hypotheses of the test ψ_{ij} and considering non-triviality, there exists a test such that $B_{ij} > 0$. Thus bounding from below by the test with maximum distance (18), $B_{\max} = \max_{\{i,j\in[\ell]:\psi_{ij}\in\tilde{\psi}\}} B_{ij} > 0$ and using the binomial expansion, we obtain (19).

Now, using Jensen's inequality, we have $B_{ij} \leq \frac{1}{2}\left(D(Q_1\|Q_2) + D(Q_2\|Q_1)\right)$. This follows from the definition of the binary hypotheses tests and the independence of samples.

From Pinsker's inequality and reverse Pinsker's inequality [23], we have

$$(2\log_2 e)\delta^2(P,Q) \leq D(P\|Q) \leq \left(\frac{4\log_2 e}{Q_{\min}}\right)\delta^2(P,Q),$$

where $Q_{\min} = \min_{x\in\text{supp}(P)} Q(x)$ and supp$(\cdot)$ is the support of the distribution. Since we are concerned with the sample complexity in the worst case when $\delta(P,Q) \to 0$, it suffices to consider $Q_{\min} > 0$. Thus, the above bounds indicate that

$$D(Q_1\|Q_2) \asymp D(Q_2\|Q_1) \asymp \theta_d^2,$$

where it is said $g \asymp h$, if there exists constants $a, b > 0$ such that $ah \leq g \leq bh$.

Thus, we have

$$P_e(\Phi) \geq \frac{1}{8}\exp\left(\log \ell - n(c\theta_d^2)\right),$$

where c is the constant scaling based on Pinsker's and reverse Pinsker's inequalities.

From this we observe that $N^*_{\text{temp}} = \Omega\left(\frac{\log \ell - \log \epsilon}{\theta_d^2}\right)$. ■

Corollary 3: Let $f(\cdot)$ be a convex function satisfying the smoothness constraints. Further, let $\min_{i,j\in[\tau]} D_f(Q_i\|Q_j) = \theta_d$. Then for constant ℓ,

$$N^*_{\text{temp}} = \Omega(\theta_d^{-1}).$$

Proof: From (10) we have

$$\min_{i,j\in[\tau]} \delta(Q_i, Q_j) \lesssim \sqrt{\theta_d}.$$

This proves the result. ■

Thus, we see that the universal clustering decoder achieves the lower bound up to the constant factor in sample complexity. Hence our clustering algorithm is asymptotically order optimal in the number of objects to be clustered and the minimum separation of hypotheses.

It is worth noting that the quantity θ_d^2 is equivalent in definition to the crowd quality defined in [5] and matches the lower bound obtained on the cost for binary classification using crowd workers biased toward giving the right label. The factor of $\log \ell$ in the cost per object arises since error probability studied here is the block (blocklength ℓ) error probability whereas [5] studies the average symbol (classification) error probability.

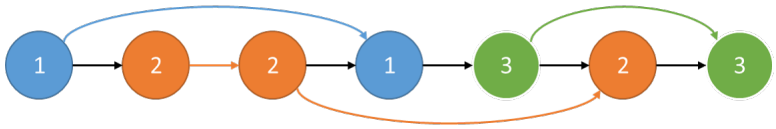

Fig. 3. Bayesian Network model of responses to a set of 7 objects chosen from a set of 3 types. We observe that the most recent response and the response to the most recent object of the same type influence every response.

IV. Workers with Memory

Recall that in Section II, we defined the structure of the stochastic kernel that determines the responses of workers with memory. In particular, we considered a Markov memory structure (2). This structure is represented in the Bayesian network depicted in Fig. 3.

Specifically, we assume that the response $Y_{i,j}$ to an object X_i by worker j is dependent on the response to the most recent object, X_{i-1}, and the response to the most recent object of the same class. This set of indices for any object X_i is given by $\mathcal{N}_i$.

A. Task Difficulty for Worker Pool

Let $\mathcal{Q}$ be the set of such Markov-structured distributions representing the worker pool. We assume the workers are chosen independently and identically from this set according to some underlying distribution. Thus, the conditional distribution of the response vector for a random worker is

$$\Pr\left[Y^\ell = y^\ell \middle| T^\ell = t^\ell\right] = \mathbb{E}\left[Q\left(Y^\ell = y^\ell \middle| T^\ell = t^\ell\right)\right] = \bar{Q}\left(y^\ell \middle| t^\ell\right),$$

where the expectation is taken over the worker distribution.

We know that a sufficient statistic of the worker responses is the empirical distribution. Asymptotically, we know that the empirical pmf converges to $\bar{Q}$ by the strong law of large numbers. It thus suffices to study the decoder with regard to this characteristic worker response pmf that retains the assumed memory properties.

Throughout the section, for any $i \in [\ell]$, denote by $\tilde{\imath}$ the index such that $\tilde{\imath} \in \mathcal{N}_i, T_i = T_{\tilde{\imath}}$.

Definition 9 (Memory Quality): The memory quality in a given pool of long-term workers is

$$\theta_m = \frac{1}{2}\left(\min_{i\in[\ell]} I(Y_i; Y_{\mathcal{N}_i}) - \max_{i\in[\ell], j<i, j\notin\mathcal{N}_i} I(Y_i; Y_{i-1}, Y_j)\right). \tag{20}$$

That is, the memory quality is defined by the difference between the information provided by the neighbors of the object and that provided by two other objects.

Since the crowd responses are defined by conditional independence across objects, we quantify the quality of the crowd in terms of the amount of memory retained by the workers. The above definition can be equivalently viewed as the *task difficulty* for a given pool of long-term crowd workers.

This informational definition finds operational significance in the coding theorems, Theorems 6 and 7.

B. Information Clustering using Neighbors

From the model of worker responses, we know that identifying the parents of each node is critical to cluster.

Lemma 2: Let $G = (V, E)$ be the Bayesian network representation of the worker responses. Let $\nu_i = \{j \in V : (j, i) \in E\}$. If $|\nu_i| = 1$, then either $T_{i-1} = T_i$ or $T_j \neq T_i$ for all $j < i$. If $|\nu_i| > 1$, $\nu_i \backslash \{i-1\}$ is in the same cluster as i.

Proof: The results follow from the model definition. ∎

We use the data processing inequality to obtain the following property that motivates the decoder construction.

Lemma 3: Let $C = \{c_1, \ldots, c_k\}$ and without loss of generality, let $c_1 < c_2 < \cdots < c_k$. Given (2), $C \in P^*$ if and only if for all $j \in [k]$,

$$c_j = \arg\max_{i<c_{j+1}} I(Y_{c_{j+1}}; Y_i, Y_{c_{j+1}-1}).$$

Algorithm 2 Clustering given MI, $(P) = \Phi_{\text{info}}(I)$

$F(i) \leftarrow 0$ for all $i \in [m]$
$P \leftarrow \emptyset$
$M_c \leftarrow 0$
$\gamma_n = c_1 n^{-\alpha}$, where c_1 is a constant and $\alpha \in (0, 1/2)$
for $i = m$ **to** 1 **do**
  if $F(i) = 0$ **then**
    $F(i) \rightarrow 1$
    $M_c \leftarrow M_c + 1$, $C(i) \leftarrow M_c$
  end if
  $I_{\max,i} \leftarrow \max_{k \leq i-1} I(Y_i; Y_{i-1}, Y_k)$
  $\eta_i = \max\{j < i : I(Y_i; Y_{i-1}, Y_j) \geq I_{\max,i} - \gamma_n\}$
  if $F(\eta_i) = 0$ **then**
    $C(\eta_i) = C(i), F(\eta_i) = 1$
  end if
end for
$P = \{\{i : C(i) = k\} : k \in [M_c]\}$

Proof: Let $i, j \in [\ell]$ such that $j < i$. Then, the result follows from the data processing inequality:

$$\begin{aligned} I(Y_i; Y_{i-1}, Y_{\tilde{\imath}}, Y_j) &= I(Y_i; Y_{\mathcal{N}_i}) + I(Y_i; Y_j | Y_{\mathcal{N}_i}) \\ &= I(Y_i; Y_{i-1}, Y_j) + I(Y_i; Y_{\tilde{\imath}} | Y_{i-1}, Y_j). \end{aligned}$$

For the given model, $I(Y_i; Y_j | Y_{\mathcal{N}_i}) = 0$. In the non-trivial problem it is natural that $I(Y_i; Y_{\tilde{\imath}} | Y_{i-1}, Y_j) > 0$. That is, the most recent object of the same class has residual information, given any other pair from the past. This in turn implies that for all $j < i, j \neq \tilde{\imath}$,

$$I(Y_i; Y_{i-1}, Y_j) < I(Y_i; Y_{\mathcal{N}_i}).$$

The result thus follows. ■

It is evident that the partition can be obtained through a careful elimination process using mutual information values. Maximum likelihood estimates of the mutual information can be obtained from the samples using and asymptotically consistent estimators. Note that such estimators converge exponentially; convergence rates are detailed in Appendix B.

C. Information Clustering Algorithm

We now describe the clustering algorithm in two stages. First we describe an algorithm that, given the set of objects and mutual information values, outputs a partition that is denser than the correct partition. We then describe an algorithm that overcomes this shortcoming by identifying sub-clusters within the identified clusters recursively. We then show correctness of the algorithm and prove it is asymptotically consistent when the ML estimates of mutual information are used.

From the directed acyclic graph (Bayesian network) corresponding to the given set of objects, we know that identifying the parents of each node is sufficient to identify clusters such that objects of the same type are in the same cluster. From Lemma 3, for any $i \in [\ell]$, $I(Y_i; Y_{i-1}, Y_j) < I(Y_i; Y_{\mathcal{N}_i})$. Thus identifying the parents of node i is equivalent to solving

$$\eta_i = \arg \max_{j \leq i-1} I(Y_i; Y_{i-1}, Y_j).$$

Using this feature we design Algorithm 2, $\Phi_{\text{info}}(I)$.

The algorithm outputs the partition of a set of objects when the corresponding mutual information values are given as input. The algorithm starts from the last object and iterates backward while finding the parents of each node. Upon identification, the parent is added to the same cluster as the object.

Theorem 5: Given a set of objects T^ℓ and the corresponding set of mutual informations $\{I(Y_i; Y_{i-1}, Y_j) : i \in [\ell], j < i\}$, the output of Alg. 2 satisfies $P = \Phi_{\text{info}}(I) \succeq P^*$.

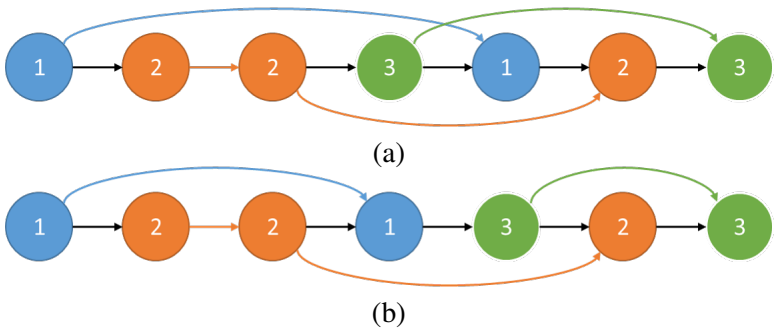

Fig. 4. Functioning and shortcoming of Alg. 2. (a) Alg. 2 outputs $P = \{\{1,5\},\{2,3,6\},\{4,7\}\} = P^*$. Here, object 4 of type 3 is not clustered with object 3 of type 2 as it is already assigned a cluster. (b) Alg. 2 outputs $P = \{\{1,4,5,7\},\{2,3,6\}\} \succeq P^*$. Here object 5 is clustered with 4 as objects of Type 1 are not encountered before.

Algorithm 3 Clustering with memory, $P = \Phi_{\text{mem}}(T^\ell)$

Choose constant $k = \lceil \frac{-\log \epsilon}{(\log \ell - 2\log \tau)} \rceil$
for $i = 1$ **to** k **do**
  Choose a permutation $\xi([\ell])$ uniformly at random
  Collect responses $\mathbf{Y}^{(i)}$ for the sequence of objects $T^{\xi([\ell])}$
  Compute $I = \{\hat{I}(Y_j; Y_{j-1}, Y_k) : k < j, j \in [\ell]\}$
  $P_i = \Phi_{\text{info}}(I)$
end for
Choose finest partition P such that $P \preceq P_i$ for all $i \in [k]$

Proof: From Lemma 3, $I(Y_i; Y_{i-1}, Y_j) \leq I(Y_i; Y_{\mathcal{N}_i})$ for all $j < i$ with equality if and only if $j \in \mathcal{N}_i$. Thus, the parents of every node $i \in [\ell]$ in the Bayesian network can be determined, given the mutual information values.

For $P = \Phi_{\text{info}}(I)$, for every object $t \in \mathcal{T}$, there exists $C \in P$ such that $\{i \in [\ell] : T_i = t\} \subseteq C$. Hence the result follows. ■

D. Consistency of Universal Clustering

Theorem 5 indicates that, given the mutual information values, the objects of the same type are clustered together. The maximizer η_i in Alg. 2 is clustered only if it is not assigned a cluster before the iteration. Thus object i is not paired with $i-1$ unless it has not been assigned a cluster. This particular scenario is depicted in the Bayesian network in Fig. 4(a).

However, the algorithm fails in a specific scenario. When there exists clusters C_1 and C_2 such that $i < j$ for every $i \in C_1, j \in C_2$, and, $\max\{i \in C_1\} = \min\{j \in C_2\} - 1$, then the resulting partition consists of the single cluster $C_1 \cup C_2$ rather than the two individual clusters. This is because objects of C_1 have not yet been encountered and thus the immediate neighbor of the first object of C_2 is clustered along with C_2 due to the Markov memory structure. This particular scenario is depicted in Fig. 4(b).

Such shortcomings of the algorithm however happen for outlier cases that are of low probability when $\ell \gg \tau$. Thus, if the finest partition in a collection of permutations of a given set of objects, chosen uniformly at random, is obtained using Φ_{info}, then with high probability, the correct partition is obtained. Thus the overall algorithm can be summarized as in Alg. 3 when we want an error probability less than 2ϵ.

For each permutation of the objects, n responses are obtained for each object from the workers. Thus, the overall number of samples per object obtained is kn.

Theorem 6: Let T^ℓ be the set of objects, $\ell \geq \tau^2$, and, let $\mathbf{Y}$ be the set of responses. Let $k \geq \lceil \frac{-\log \epsilon}{(\log \ell - 2\log \tau)} \rceil$ be the number of permutations chosen in Alg. 3. Then, for

$$n \gtrsim \max\left\{ (\log \ell - \log \epsilon)^{\frac{1}{(1-2\alpha-\beta)}}, (\log \ell - \log \epsilon)^{\frac{1}{(1-4\alpha)}}, \theta_m^{\frac{-1}{\alpha}} \right\}, \tag{21}$$

for $0 < \alpha < 1/2$ and $0 < \beta < 1$ such that $\log_2^2 n \leq n^\beta$, $P_e(\Phi_{\text{mem}}) \leq 2\epsilon$, for any $\epsilon > 0$. Further, for constant ℓ and θ_m, the algorithm is asymptotically consistent.

Proof: We first observe that when $|\hat{I}(Y_i;Y_{i-1},Y_j) - I(Y_i;Y_{i-1},Y_j)| < \theta_m$ for all $i \in [\ell], j < i$, $\Phi_{\text{info}}(\hat{I}) = \Phi_{\text{info}}(I)$. That is, when the empirical mutual information values do not deviate from the actual values significantly, the clustering algorithm works without error.

For $n \geq (c_1/\theta_m)^{1/\alpha}$, $\gamma_n < \theta_m$. Let $I_{ij} = I(Y_i;Y_{i-1},Y_j)$. Then,

$$\begin{aligned}
\Pr\left[\Phi_{\text{info}}(\hat{I}) \neq \Phi_{\text{info}}(I)\right] &\leq \sum_{i\in\ell, j<i} \Pr\left[|\hat{I}_{ij} - I_{ij}| > \gamma_n\right] && (22)\\
&\leq \binom{\ell}{2} \exp\left(\frac{-n\gamma_n^2}{18\log_2^2 n} + o(1)\right) && (23)\\
&\leq \exp\left(2\log\ell - \frac{c_1^2}{18} n^{(1-\nu)} + o(1)\right),
\end{aligned}$$

implying asymptotic consistency. Here $\nu = 2\alpha + \beta$, (22) follows from the union bound, and (23) follows from Lemma 6.

Additionally, we note that

$$\begin{aligned}
\Pr\left[\Phi_{\text{info}}(\hat{I}) \neq \Phi_{\text{info}}(I)\right] &\leq \sum_{i\in\ell, j<i} \Pr\left[|\hat{I}_{ij} - I_{ij}| > \gamma_n\right]\\
&\leq 3\binom{\ell}{2} (n+1)^{\tau^2} \exp\left(-\tilde{c} n \gamma_n^4\right) && (24)\\
&\leq \exp\left(2\log\ell - \tilde{c} c_1^4 n^{(1-4\alpha)} + o(1)\right),
\end{aligned}$$

where (24) follows from (47).

To obtain the correct partition, we use the responses generated for several uniformly random permutations of the given set of objects and select the finest partition. The correct partition may not be recovered when there exists $t, t' \in \mathcal{T}$, such that $\max\{i \in [\ell] : T_i = t\} = \min\{i \in [\ell] : T_i = t'\} - 1$ in every chosen partition.

Let $K_t = |\{i \in [\ell] : T_i = t\}|$, $t \in [\tau]$ and let $M_t = \max\{i \in [\ell] : T_i = t\}$ and $m_t = \min\{i \in [\ell] : T_i = t\}$. Thus, the probability that $P = \Phi_{\text{mem}}(T^\ell) \succ P^*$ is bounded as:

$$\begin{aligned}
\Pr\left[P \succ P^*\right] &= \Pr\left[\exists t \neq t' : M_t = m_{t'} - 1\right]^k\\
&\leq \left(\sum_{t,t'\in\mathcal{T}, t\neq t'} \Pr\left[M_t = m_{t'} - 1\right]\right)^k.
\end{aligned}$$

The total number of possible permutations of the given set of objects is given by

$$\kappa_{\text{permut}} = \frac{\ell!}{\prod_{\tilde{t}\in\mathcal{T}} K_{\tilde{t}}!}.$$

The number of sequences such that $M_t = m_{t'} - 1$ can be determined by choosing $K_t + K_{t'} - 1$ locations out of $\ell - 1$ locations to fill the objects of type t and t' and permute over the other objects. Thus,

$$\kappa_{\text{t,t'}} \leq \binom{\ell - 1}{K_t + K_{t'} - 1} \frac{(\ell - K_t - K_{t'})!}{\prod_{\tilde{t}\in\mathcal{T}\setminus\{t,t'\}} K_{\tilde{t}}!}.$$

Since the permutations are chosen uniformly at random,

$$\begin{aligned}
\Pr\left[m_{t'} - M_t = 1\right] &\leq \frac{\kappa_{t,t'}}{\kappa_{\text{permut}}}\\
&= \frac{1}{\ell} \frac{K_t! K_{t'}!}{(K_t + K_{t'} - 1)!} \leq \frac{1}{\ell}, && (25)
\end{aligned}$$

where (25) follows from the fact that

$$\frac{K_t! K_{t'}!}{(K_t + K_{t'} - 1)!} \leq 1, \tag{26}$$

as $K_t, K_{t'} \geq 1$.

Thus, we have

$$\Pr\left[P \succ P^*\right] \leq \left(\binom{\tau}{2}\frac{2}{\ell}\right)^k \leq \exp\left(-k(\log \ell - 2\log \tau)\right). \tag{27}$$

Thus, for $k \geq \lceil \frac{-\log \epsilon}{(\log \ell - 2\log \tau)} \rceil$, $\Pr\left[P \succ P^*\right] \leq \epsilon$.

Now we prove consistency of Alg. 3. Using the union bound, we observe that the probability of error is bounded as

$$\begin{aligned} P_e &\leq \Pr\left[P \succ P^*\right] + k \Pr\left[\Phi_{\text{info}}(\hat{I}) \neq \Phi_{\text{info}}(I)\right] \\ &\leq \exp\left(-k(\log \ell - 2\log \tau + \log 2)\right) \\ &\quad + \exp\left(2\log \ell - \max\left\{\frac{c_1^2}{18} n^{(1-\nu)}, \tilde{c} c_1^4 n^{(1-4\alpha)}\right\} + o(1)\right) \\ &\leq 2\epsilon \end{aligned} \tag{28}$$

for a large enough n. Here (28) follows from the two concentration bounds on empirical mutual information described in Appendix B.

Thus, for any $\epsilon > 0$, there exists n, k sufficiently large, such that $P_e < \epsilon$. Hence the consistency result follows. The sample complexity is obtained from the error exponent in (28). ∎

We observe from the proof that there is a trade-off between the values of k and n needed to achieve a certain level of accuracy. In particular, we observe that when ℓ is large, it suffices to consider a small number of permutations of the set of objects, while each permutation requires a larger number of samples. On the other hand, when ℓ is relatively small, one needs a large number of permutations while each permutation requires far fewer samples.

We restrict focus to the case where $\ell \gg \tau^2$ and under this scenario find the following result on sample complexity.

Corollary 4: Given $\mathcal{T} = [\tau]$ with $\tau < \infty$ a constant and $\ell \gg \tau^2$,

$$N^*_{\text{mem}}(\epsilon) = O\left(\frac{(\log \ell - \log \epsilon)^{\min\{1/(1-2\alpha-\beta), 1/(1-4\alpha)\}}}{\theta_m^{(2/(1-\beta))}}\right).$$

Proof: Using Theorem 6 and the fact that the total number of samples used is kn (since clustering with n samples is done k times) per object, we obtain the result. ∎

For large ℓ, we can thus observe that $N^*_{\text{mem}}(\epsilon) = O\left(\frac{\log \ell}{\theta_m^2}\right)$.

Note that under the Markov memory model for long-time workers, the sufficient number of samples per object is the same in order as for temporary workers.

E. Lower Bound on Sample Complexity

We now provide matching lower bounds by studying the probability of error of a problem which is a reduction of the universal clustering problem.

Theorem 7: The sample complexity of universal clustering using workers with memory satisfies

1) for a fixed $\theta_m > 0$, $N^*_{\text{mem}} = \Omega\left(\log \ell\right)$, and
2) for a fixed $\ell < \infty$, $N^*_{\text{mem}} = \Omega\left(\theta_m^{-1}\right)$.

Proof: Choose a prior, parametrized by the size of the problem ℓ as $P_T(1) = 1 - \frac{1}{\ell}, P_T(2) = \frac{1}{\ell}$ such that $\frac{1}{\ell} = \frac{1}{\ell}$. Let $\mathcal{E}$ be the set of all vectors of objects with at most one object of type 2. Then,

$$\Pr\left[T^\ell \in \mathcal{E}\right] = \left(2 - \frac{1}{\ell}\right)\left(1 - \frac{1}{\ell}\right)^{\ell-1} \overset{\ell \to \infty}{\longrightarrow} 1.$$

In particular, we note that $\Pr\left[T^\ell \in \mathcal{E}\right]$ is an increasing function of ℓ and is at least $\frac{1}{2}$ for any $\ell > 3$.

For a given constant θ_m, consider the special case of the problem where $\mathcal{N}_i = \{\tilde{\imath}\}$. That is, consider the problem where any two objects are dependent if and only if they are of the same type. Clearly, any algorithm that solves the

universal clustering with memory problem solves this simplified problem as well. Thus, following the convention established, we have

$$\begin{cases} I(Y_i;Y_{\tilde{i}}) \geq 2\theta_m, & \text{for all } i \in [\ell] \\ I(Y_i;Y_j) = 0, & \text{for all } i,j \in [\ell], T_i \neq T_j \end{cases}.$$

Define

$$W = [W_{ij}]_{1\leq i,j\leq 2} = \begin{bmatrix} \frac{1}{2}+\epsilon & \frac{1}{2}-\epsilon \\ \frac{1}{2}-\epsilon & \frac{1}{2}+\epsilon \end{bmatrix}.$$

Consider the scenario where worker responses are inertial over time and characterized as:

$$\Pr\left[Y_i = k | Y_{\tilde{i}} = j\right] = W_{kj},$$

for any $k, j \in \{0,1\}$ and $i \in [\ell]$. Additionally, assume that the marginals of the responses are uniform (that is, the response to the first object of each type is distributed as Bern$(1/2)$).

The information constraint implies

$$\frac{1}{2} - h^{-1}(1-2\theta_m) \leq \epsilon < \frac{1}{2},$$

where $h(\cdot)$ is the binary entropy function and $h^{-1}(\cdot)$ is its inverse. Let $\epsilon = \frac{1}{2} - h^{-1}(1-2\theta_m)$.

From the definition of the error probability,

$$P_e \geq \frac{1}{2}\Pr\left[\hat{P} \neq P^* | T^\ell \in \mathcal{E}\right]. \tag{29}$$

Now consider the set $\mathcal{E}$ of vectors. Identifying the correct partition for a vector of objects from this space is equivalent to identifying the objects. Thus consider the $(\ell+1)$-ary hypothesis testing problem defined by

$$\begin{cases} H_0 : T_j = 1, \text{ for all } j \in [\ell] \\ H_i : T_i = 2, T_j = 1, \text{ for all } j \neq i \end{cases}. \tag{30}$$

We seek to compute the average error probability of (30) corresponding to the prior P_T. Due to symmetry, note that the optimal decoder accrues the same probability of error under H_i for any $i > 0$. Thus

$$\begin{aligned} \Pr\left[\text{error in (30)}\right] = \Pr\left[H_0\right] \Pr\left[\text{error in (30)} | H_0\right] \\ + \Pr\left[\text{error in (30)} | H_1\right] \left(\sum_{i\in[\ell]} \Pr\left[H_i\right]\right). \end{aligned}$$

Now, note that

$$\sum_{i\in[\ell]} \Pr\left[H_i\right] = \left(1-\frac{1}{\ell}\right)^{\ell-1}, \quad \Pr\left[H_0\right] = \left(1-\frac{1}{\ell}\right)^{\ell}.$$

Thus, for $\ell > 1$,

$$\frac{1}{2}\Pr\left[\{H_i : i > 0\}\right] \leq \Pr\left[H_0\right] \leq \Pr\left[\{H_i : i > 0\}\right].$$

Thus, $\Pr\left[\{H_i : i > 0\}\right] \asymp \Pr\left[H_0\right]$. This indicates that the average error probability is lower-bounded by a constant factor of the minimax error probability lower bound for (30).

Let Q_i be the distribution of the set of responses corresponding to the hypotheses defined in (30):

$$Q_i(Y^\ell) = \begin{cases} \frac{1}{2}\prod_{k=2}^{\ell} W_{Y_k Y_{k-1}}, & i = 0 \\ \frac{1}{4}\prod_{j=2}^{i-1} W_{Y_j Y_{j-1}} \prod_{k=i+1}^{\ell} W_{Y_k Y_{k-1}}, & i \in [\ell] \end{cases}. \tag{31}$$

Lemma 4: For all ℓ, $D\left(Q_i \| Q_j\right) = O(1)$.

Proof: See Appendix C. ∎

Having bounded the KL divergences between the hypotheses, we obtain a lower bound on the error probability of (30) using the generalized Fano inequality [24].

Let $\beta = \max_{i,j\in[\ell]\cup\{0\},i\neq j} D(Q_i\|Q_j)$. The loss function considered here is the 0-1 loss. Hence,

$$\begin{aligned}\Pr\left[\hat{P} \neq P^*|T^\ell \in \mathcal{E}\right] &= \Pr\left[\text{error in problem (30)}\right] \\ &\asymp \max_{0\leq i\leq \ell} \Pr\left[\text{error in problem (30)}|H_i\right] \\ &\geq \frac{1}{2}\left(1 - \frac{n\beta + \log 2}{\log(\ell+1)}\right). \end{aligned} \tag{32}$$

Hence, for a constant $\theta_m > 0$, the sample complexity of universal clustering satisfies

$$N^*_{\text{mem}} = \Omega(\log \ell).$$

Now, when ℓ is fixed, we seek to understand the sample complexity with respect to the memory quality of the crowd. To this end, we note that any consistent clustering algorithm is also consistent for the binary hypothesis test

$$\psi : \begin{cases} H_0 : I(Y_1;Y_2) = 0 \\ H_1 : I(Y_1;Y_2) \geq 2\theta_m \end{cases}. \tag{33}$$

That is, if Φ is a decoder for the universal clustering problem, then Φ also solves ψ.

Since the sufficient statistics for detection of the binary hypothesis testing above are the responses to T_1, T_2, it suffices to consider Y_1^n, Y_2^n. Let the prior here be $P_T(1) = P_T(2) = 1/2$. Let the corresponding distributions of worker responses be $p(Y_1^n, Y_2^n)$ and $q(Y_1^n, Y_2^n)$ under H_0 and H_1 respectively. Here,

$$\begin{aligned} p(y_1^n, y_2^n) = &\frac{1}{2}\prod_{i=1}^n \Pr\left[y_{1,i}, y_{2,i}|T_1 = T_2 = 1\right] \\ &+ \frac{1}{2}\prod_{i=1}^n \Pr\left[Y_{1,i}, Y_{2,i}|T_1 = T_2 = 2\right], \end{aligned}$$

and

$$\begin{aligned} q(Y_1^n, Y_2^n) = &\frac{1}{2}\prod_{i=1}^n \Pr\left[Y_{1,i}|T_1 = 1\right]\Pr\left[Y_{2,i}|T_1 = 2\right] \\ &+ \frac{1}{2}\prod_{i=1}^n \Pr\left[Y_{1,i}|T_1 = 2\right]\Pr\left[Y_{2,i}|T_1 = 1\right] \end{aligned}$$

Let $p_j^{(i)}(y) = \Pr\left[Y_i = y|T_i = j\right]$, $q_j^{(i)}(y) = \Pr\left[Y_i = y|T_i = j\right]$, under H_0 and H_1 respectively. Without loss of generality, we assume $I_p(Y_1, Y_2) = 2\theta_m$.

Since the distributions satisfy the information constraints, we have,

$$\begin{aligned} D(p\|q) &\leq \frac{n}{4}\left(4I_p(Y_1;Y_2) + \sum_{i,j,k\in[2]} D(p_j^{(i)}\|q_k^{(i)})\right) \\ &= 2n\theta_m, \end{aligned} \tag{34}$$

when the marginals under the two hypotheses are equal as was the case in the inertial worker response channel. Here (34) follows from convexity. Thus the minimum upper bound on the KL divergence between the hypotheses is $2\theta_m$. Since we consider the worst case with respect to θ_m, it suffices to consider this upper bound.

Thus,

$$P_e(\Phi) \geq \Pr\left[\text{error in } \psi\right] \geq \frac{1}{4}\exp\left(-2B(p,q)\right) \tag{35}$$

$$\geq \frac{1}{4}\exp\left(-2D(p\|q)\right) \geq \frac{1}{4}\exp\left(-4n\theta_m\right), \tag{36}$$

where (35) follows from the Kailath lower bound [22]. Then, using Jensen's inequality, we obtain (36). Thus

$$N^*_{\text{mem}} = \Omega\left(\theta_m^{-1}\right).$$

■

We observe from Theorem 7 that the universal clustering algorithm is order optimal in terms of the number of objects, ℓ. However, there is a gap between the lower bound and the achievable cost in terms of the crowd quality θ_m. This gap is exactly the well-known gap for entropy estimation observed in [25, Corollary 2].

F. Reductions to Other Clustering Algorithms

There exist several clustering paradigms based on mutual information. Here we describe two such algorithms and reductions of our model that lead to our clustering algorithm becoming the same as those algorithms.

An information clustering strategy is defined in [14] that identifies clusters based on the minimum partition information that is defined as

$$I(Z_V) = \min_{P \in \mathcal{P}} I_P(Z_V),$$

where I_P is the partition information according to the partition P, defined as

$$I_P(Z_V) = \frac{1}{|P| - 1} \left(\sum_{C \in P} H(Z_C) - H(Z_V) \right).$$

Consider the Markov memory model such that $\mathcal{N}_i = \{\tilde{i}\}$, that is, if Y_i is conditionally dependent on an object only if it is of the same type. Then, if $|P^*| > 1$, then

$$I(Y^\ell) = I_{P^*}(Y^\ell) = 0.$$

That is, the correct partition is the partition that minimizes the partition information. The following reduction indicates that minimizing the partition information is the same as our algorithm, and so [14] is a special case of our approach.

First, we have

$$H(Y^\ell) = \sum_{i=1}^{\ell} H(Y_i | Y_{\tilde{i}}).$$

Similarly,

$$H(Y_C) = \sum_{i \in C} H(Y_i | Y_{j(i)}),$$

where $j(i) = \max\{k < i : k \in C\}$. This implies that

$$I_P(Y^\ell) = \frac{1}{|P| - 1} \left(\sum_{i=1}^{\ell} (I(Y_i; Y_{\tilde{i}}) - I(Y_i; Y_{j(i)})) \right),$$

where $j(i) = \max\{k < i : i, k \in C\}$. This indicates that minimizing the partition information is equivalent to $\Phi_{\text{info}}(I)$.

Another information-based clustering strategy is the mutual information relevance network (MIRN) clustering [17]. Here, for a given threshold γ, the clustering strategy determines the connected components of $G = ([\ell], E)$ such that $(i, j) \in E \iff I(Z_i; Z_j) > \gamma$. For the Markov memory model, under the restriction that

$$\min_{\{i,j \in [\ell] : T_i = T_j\}} I(Y_i; Y_j) > \gamma \geq \max_{\{i,j \in [\ell] : T_i \neq T_j\}} I(Y_i; Y_j),$$

MIRN outputs the correct partition.

However, in the universal clustering scenario, the decoder is not aware of γ and thus, it may not be feasible to implement MIRN clustering optimally. Nevertheless. under the restriction that $\mathcal{N}_i = \{\tilde{i}\}$ and $\gamma = 0$, the MIRN clustering algorithm is just a special case of $\Phi_{\text{info}}(I)$.

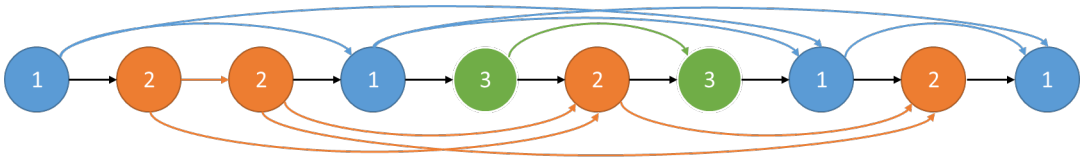

Fig. 5. Bayesian Network model of responses of to a set of 10 objects chosen from a set of 3 types with $\zeta = 2$. We observe that the most recent response and the response to the two most recent object of the same type influence every response.

G. Extended Memory Workers

While the model defined above considers the dependence of responses on just the most recent object of the same kind, our results hold for any fixed, finite-order Markov structure as well. In particular, consider the scenario where worker responses are dependent on the set $\mathcal{N}_i = \{i-1\} \cup \mathcal{M}_i$, where $\mathcal{M}_i \subseteq \{j < i : T_j = T_i\}$ such that it contains the most recent ζ indices of the same type of object. That is, the response to an object is dependent not only on the most recent response, but also a constant number of prior responses to objects of the same type.

An example of this worker model can be found in Fig. 5 for a set of 10 objects of 3 types with a worker memory of $\zeta = 2$.

Then the algorithm defined and the results obtained can be extended to this scenario. More specifically, the parents of a node i can be computed using the rule

$$I(Y_i; Y_{\mathcal{N}_i}) > I(Y_i; Y_S),$$

for any index set $S \subseteq [i-1]$ such that $|S| \leq \zeta$.

Then for any constant ζ and $n \gtrsim (\tau+1)^{2\zeta}$, the consistency of the algorithm holds. That is, as long as $\zeta = O(1)$, the sample complexity results follow.

V. Unified Worker Model

While we studied two distinct classes of worker models in temporary workers and workers with memory, these two scenarios are limiting cases of a unified worker model described here. After all, it is reasonable to characterize practical crowd worker decisions as influenced by both aspects—memory of individual responses and task difficulty with respect to objects.

For the unified model, we provide an achievable scheme that makes use of the algorithms defined earlier. Further, we prove consistency and order optimality of the scheme. As in Section II, consider worker model (3), where each worker is characterized by a Markov memory model subject to fixed conditional marginal distributions.

A. Worker Quality

Let us now define worker quality. Let $Q_1, \ldots, Q_\tau$ be the conditional response distributions given the object class, defined as

$$Q_i(y) \triangleq \Pr[Y = y | T = i] = \mathbb{E}\left[Q^{(W)}(Y = y | T = i)\right],$$

where the expectation is taken over the workers in the pool. Define *distance quality* as

$$\theta_d \triangleq \min_{\{i,j \in \mathcal{T}, i \neq j\}} \delta(Q_i \| Q_j).$$

Notice this *distance quality* is analogous to the definition of the worker quality in the case of temporary workers.

Additionally, define the *memory quality* as

$$\theta_m \triangleq \frac{1}{2}\left(\min_{i \in [\ell]} I(Y_i; Y_{\mathcal{N}_i}) - \max_{i \in [\ell], j < i, j \notin \mathcal{N}_i} I(Y_i; Y_{i-1}, Y_j)\right).$$

Notice this *memory quality* is analogous to the definition of worker quality defined for the worker with memory scenario.

Worker quality in the unified worker model is a combination of these individual quality parameters.

Algorithm 4 Clustering under unified worker model, $\Phi_{\text{u}}(T^\ell)$

$P_{\text{info}} \leftarrow \Phi_{\text{mem}}(T^\ell)$
for $C \in P_{\text{info}}$ **do**
  $P_C \leftarrow \Phi_{\text{temp}}(T^C)$
end for
$P \leftarrow \cup_{C \in P_{\text{info}}} P_C$

B. Clustering Algorithm

We now provide the universal clustering strategy for the unified worker model, Alg. 4, drawing on achievable schemes from before. First perform the memory-based clustering defined in Alg. 3; then for every cluster in the partition output by the algorithm, perform distance-based clustering. We now show the consistency of the algorithm.

Theorem 8: Let T^ℓ be the set of objects and $\ell > \tau^2$. Then, for

$$n \gtrsim \max\Big\{ (\log \ell - \log \epsilon)^{\frac{1}{(1-2\alpha-\beta)}}, \tag{37}$$

$$(\log \ell - \log \epsilon)^{\frac{1}{(1-4\alpha)}}, (\theta_m + \theta_d)^{\frac{-1}{\alpha}} \Big\}, \tag{38}$$

for $0 < \alpha < 1/2$ and $0 < \beta < 1$, $P_e(\Phi_{\text{u}}) \leq 2\epsilon$, for any $\epsilon > 0$.

Proof: First we note that for $n \geq (c_1/(\theta_m + \theta_d/4))^{1/\alpha}$, $\gamma_n \leq \theta_d/4 + \theta_m$. Thus, at least one of $\gamma_n \leq \theta_d/4$ or $\gamma_n \leq \theta_m$ is true. This in turn indicates that at least one of Φ_{mem} or Φ_{temp} is consistent.

Next, from Theorem 5, we note that the output $P_{\text{info}} \succeq P^*$. Thus, subsequent clustering of the individual clusters is sufficient. This in turn indicates the correctness and asymptotic consistency of Alg. 4. ∎

We now observe that the sample complexity with respect to the number of objects to be clustered is still $O(\log \ell)$ while that with respect to the quality parameters is $O((\theta_m + \theta_d)^{-2})$.

Corollary 5: Given $\mathcal{T} = [\tau]$ with $\tau < \infty$ a constant,

1) for a constant $\theta_m, \theta_d > 0$, $N^*_{\text{u}}(\epsilon) = O(\log \ell)$;
2) for a constant ℓ, $N^*_{\text{u}}(\epsilon) = O((\theta_d + \theta_m)^{-2})$.

It is worth noting the limiting cases of the unified worker model. In particular, when $\theta_m \to 0$, the problem reduces to clustering with temporary workers as do the achievable scheme and sample complexity requirements. On the other hand, $\theta_d \to 0$ corresponds to a particular case of clustering using workers with memory.

C. Lower Bound on Sample Complexity

We now derive the lower bound on sample complexity by extending the proof of the converse for workers with memory.

Theorem 9: Sample complexity of universal clustering under the unified worker model satisfies

1) for a fixed $\theta > 0$, $N^*_u = \Omega(\log \ell)$, and
2) for a fixed $\ell < \infty$, $N^*_u = \Omega\left((\theta_m + \theta_d^2)^{-1}\right)$.

Proof: We proceed in similar fashion to the proof for the case of workers with memory. Again, consider the prior parametrized by the size of the problem ℓ as $P_T(1) = 1 - \pi, P_T(2) = \pi$ such that $\pi = \frac{1}{\ell}$. Again, we will use the generalized Fano's inequality over the space $\mathcal{E}$ of vectors of objects. We again consider the case of $\mathcal{N}_i = \{\tilde{i}\}$.

Consider worker responses such that marginals of the responses to an object satisfy

$$\Pr[Y = i | T = j] = \begin{cases} p, & i = j \\ 1 - p, & i \neq j \end{cases}$$

irrespective of the order of occurrence. Define the matrices

$$W^{(1)} = [W^{(1)}_{ij}]_{1 \leq i,j \leq 2} = \begin{bmatrix} a & 1-a \\ 1-b & b \end{bmatrix},$$

and

$$W^{(2)} = [W^{(2)}_{ij}]_{1 \leq i,j \leq 2} = \begin{bmatrix} b & 1-b \\ 1-a & a \end{bmatrix}.$$

Let the worker responses be characterized by

$$\Pr\left[Y_i = k \middle| Y_{\tilde{i}} = \tilde{k}, T_i = T_{\tilde{i}} = j\right] = W_{k\tilde{k}}^{(j)}.$$

From the constraint on distance quality, we have:

$$2p - 1 \geq \theta_d.$$

The constraint on the nature of the marginals establishes:

$$ap - b(1-p) = p.$$

The restriction on the information quality implies:

$$h(p) - ph(a) - (1-p)h(b) \geq 2\theta_m.$$

Let us consider the case when both inequalities hold with equality. This yields a specific worker channel that satisfies the memory and distance quality requirements. We analyze the error probability on this worker channel.

Again, using analysis similar to the proof of Lemma 4, we observe the KL divergences between the hypotheses in the $(\ell+1)$-ary hypothesis testing problem are $O(1)$. Hence there exists a constant β such that (32) holds. Thus, for constant θ_m and θ_d, the sample complexity of universal clustering satisfies:

$$N_u^* = O(\log \ell).$$

Now, when ℓ is fixed, we study the necessary sample complexity of universal clustering with respect to θ_m, θ_d. We know that a consistent universal clustering algorithm also solves the binary hypothesis test

$$\psi : \begin{cases} H_0 : I(Y_1; Y_2) = 0 \\ H_1 : I(Y_1; Y_2) \geq 2\theta_m. \end{cases}$$

Following the analysis from the proof of Theorem 7, from (34), we have

$$D(p\|q) \leq n\left(2\theta_m + \theta_d \log\left(\frac{1+\theta_d}{1-\theta_d}\right)\right) \lesssim n(\theta_m + \theta_d^2).$$

Finally, using the Kailath lower bound, we obtain

$$P_e(\Phi) \geq \frac{1}{4}\exp\left(-4cn(\theta_m + \theta_d^2)\right). \tag{39}$$

Thus, for constant ℓ,

$$N_u^* = \Omega((\theta_m + \theta_d^2)^{-1}).$$

■

From the theorem, we note the universal clustering algorithm is order optimal in sample complexity in terms of the number of objects, for a crowd of given quality. However, for a given number of objects, there exists an order gap between achievable sample complexity and the converse. As expected, the gap follows from the gap in the case of workers with memory, which in turn is from the gap in estimating entropy [25].

In particular, we observe that in the limit of $\theta_d \to 0$, the problem reduces to the case of workers with memory and on the other hand the case of $\theta_m \to 0$ reduces to the problem of clustering using temporary workers without memory.

A finer point in the analysis to be noted is that the worst-case channels considered in the converse proofs is the inertial channel considered in the proof of Theorem 7, which is also the solution to the set of constraints for the channel in the unified scenario under the limit of $\theta_d \to 0$.

Thus, we observe temporary workers and long-term workers with memory are indeed closely related through the unified worker model studied here and are limiting scenarios.

VI. Conclusion

This paper establishes an information-theoretic framework to study the universal crowdsourcing problem. Specifically, we defined a unified worker model (incorporating aspects of human decision making from experimental crowdsourcing and behavioral economics) and designed unsupervised clustering algorithms that are budget optimal. We first studied two limiting cases of workers—ones with and ones without memory.

For temporary workers without memory, we used distributional identicality of responses to design a universal clustering algorithm that is asymptotically consistent and order optimal in sample complexity. For workers with memory, under a Markov model of memory, we used the dependence structure to design a novel universal clustering algorithm that is asymptotically consistent and order optimal in sample complexity with respect to the number of objects. We also note that the gap obtained between necessary and sufficient conditions on sample complexity with respect to the memory quality is also observed in the empirical estimation of entropy.

We then integrated the limiting cases to develop a universal clustering algorithm for the unified worker model. We again proved asymptotic consistency and order optimality in sample complexity with respect to the number of objects. With regard to the quality of crowd workers, the gap observed in the case of the workers with memory remains.

Behavioral experiments using crowd workers on platforms such as Amazon MTurk can be performed to gain insight into the performance of the algorithms in practice and to validate the unified worker models.

Our results provide a way to compare costs between crowdsourcing platforms having workers with and without memory, thereby providing the opportunity to choose the right task-dependent platform. Further, they provide a window into more general studies of the computational capabilities and complexities of human-based information systems. In particular, the work sheds light on the influence of various attributes of crowd workers such as object-specific memory. In essence, the work studies a space-time tradeoff for human computation systems and to the best of our knowledge is the first of its kind.

Appendix A
Concentration of Empirical Distributions

In this section we briefly study the rates of convergence of the ML estimates of f-divergence.

Let $\mathcal{Z}$ be a finite set of objects and $Z_1, \ldots, Z_n \overset{iid}{\sim} p$. Let $\hat{p}$ be the empirical distribution obtained as

$$\hat{p}(z) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}\{Z_i = z\}.$$

Lemma 5: If p and $\hat{p}$ are the true and empirical distributions respectively, then

$$\Pr\left[\delta(\hat{p}, p) \geq \epsilon\right] \leq (n+1)^{|\mathcal{Z}|} \exp\left(-c_0 n \epsilon^2\right), \tag{40}$$

where $c_0 = 2 \log_2 e$. Further, for any convex function f satisfying the smoothness constraints,

$$\Pr\left[D_f(\hat{p} \| p) \geq \epsilon\right] \leq (n+1)^{|\mathcal{Z}|} \exp\left(-n\epsilon / C\right), \tag{41}$$

where $C < \infty$ is a constant such that $x f''(x) < C$.

Proof: The results follow from Pinsker's inequality (9), Theorem 1, and Sanov's theorem. ∎

Appendix B
Estimating Mutual Information from Samples

Here we briefly describe the maximum likelihood (ML) estimate of mutual information and its convergence properties.

Let $X \sim p$ be a random variable on a discrete space $\mathcal{X}$. Let $X_1, \ldots, X_n \overset{\text{iid}}{\sim} p$ and let $\hat{p}$ be the corresponding empirical distribution. The ML estimate of entropy of X is given by

$$\hat{H}(X) = \mathbb{E}_{\hat{p}}\left[-\log(\hat{p}(X))\right].$$

The ML estimate of mutual information between random variables X, Y is then given by

$$\hat{I}(X;Y) = \hat{H}(X) + \hat{H}(Y) - \hat{H}(X,Y).$$

The ML estimates of entropy and mutual information have been widely studied [25]–[28]. In particular, the following results are notable:

1) For all n, from Jensen's inequality, $\mathbb{E}[\hat{H}(X)] < H(X)$ [26]. For all n, $\epsilon > 0$,
$$\Pr\left[|\hat{H}(X) - \mathbb{E}[\hat{H}(X)]| > \epsilon\right] \leq 2\exp\left(\frac{-n\epsilon^2}{2\log_2^2 n}\right), \tag{42}$$
from McDiarmid's inequality.
2) ML estimate, $\hat{H}$ is negatively biased [27]:
$$b_n(\hat{H}) \triangleq \mathbb{E}_p\left[\hat{H}(X)\right] - H(X) = -\mathbb{E}_p\left[D(\hat{p}\|p)\right] < 0.$$
Further,
$$-\frac{|\mathcal{X}|}{n} \leq -\log\left(1 + \frac{|\mathcal{X}|-1}{n}\right) \leq b_n(\hat{H}) \leq 0.$$
3) From [25], we know that the lower bound on sample complexity for estimating entropy up to an additive error of ϵ is $\frac{|\mathcal{X}|}{\epsilon \log |\mathcal{X}|}$.
4) From [28], we observe that the deviation of the empirical entropy from the entropy of $X \sim P$, is bounded in terms of the total variational distance as
$$|\hat{H}(X) - H(X)| \leq -2\delta(\hat{P}, P)\log\frac{2\delta(P,Q)}{|\mathcal{X}|}. \tag{43}$$

Since we deal with finite, constant alphabet sizes, it suffices for us to consider the ML estimates with sufficiently large n, such that the bias is negligible.

Lemma 6: For fixed alphabet sizes, $|\mathcal{X}|, |\mathcal{Y}|$, the ML estimate of entropy and mutual information are asymptotically consistent and satisfy
$$\Pr\left[|\hat{H}(X) - H(X)| > \epsilon\right] \leq 2\exp\left(\frac{-n\epsilon^2}{2\log_2^2 n} + o(1)\right) \tag{44}$$
$$\Pr\left[|\hat{I}(X;Y) - I(X;Y)| > \epsilon\right] \leq 6\exp\left(\frac{-n\epsilon^2}{18\log_2^2 n} + o(1)\right). \tag{45}$$

Proof: The convergence of entropy follows directly by applying the triangle inequality, union bound, and (42). The result follows from the fact that the alphabet is of finite, constant size. This implies the convergence result for mutual information. ∎

Lemma 7: For fixed alphabet sizes, $|\mathcal{X}|, |\mathcal{Y}|$, the ML estimate of entropy and mutual information are asymptotically consistent and satisfy
$$\Pr\left[|\hat{H}(X) - H(X)| > \epsilon\right] \leq (n+1)^{|\mathcal{X}|}\exp\left(-cn\epsilon^4\right), \tag{46}$$
$$\Pr\left[|\hat{I}(X;Y) - I(X;Y)| > \epsilon\right] \leq 3(n+1)^{|\mathcal{X}||\mathcal{Y}|}\exp\left(-\tilde{c}n\epsilon^4\right), \tag{47}$$
where $c = \left(2|\mathcal{X}|^2\log 2\right)^{-1}$, $\tilde{c} = \left(32\max\{|\mathcal{X}|,|\mathcal{Y}|\}^2\log 2\right)^{-1}$.

Proof: We first observe that for all $x > 0$, $\log x < \sqrt{x}$. Thus, for any $X \sim P$, from (43), we have
$$|\hat{H}(X) - H(X)| \leq \sqrt{2|\mathcal{X}|\delta(\hat{P}, P)}.$$
Using this and (40), the first inequality is obtained. Subsequently, using the triangle inequality and union bound, we obtain the convergence of the empirical mutual information. ∎

These rates of convergence are used to prove consistency.

APPENDIX C
PROOF OF LEMMA 4

In this section we describe the proof of Lemma 4.

Proof: We first note that

$$\begin{aligned}
D(Q_i\|Q_j) &= \mathbb{E}_{Q_i}\left[\log\left(\frac{Q_i(Y^\ell)}{Q_j(Y^\ell)}\right)\right] \\
&= \mathbb{E}_{Q_i}\left[\log\left(\frac{W_{Y_{i+1}Y_{i-1}}W_{Y_jY_{j-1}}W_{Y_{j+1}Y_j}}{W_{Y_{i+1}Y_i}W_{Y_iY_{i-1}}W_{Y_{j+1}Y_{j-1}}}\right)\right] \\
&= \mathbb{E}_{Q_{\ell-i}}\left[\log\left(\frac{W_{Y_{\ell-i+1}Y_{\ell-i-1}}W_{Y_{\ell-j}Y_{\ell-j-1}}W_{Y_{\ell-j+1}Y_{\ell-j}}}{W_{Y_{\ell-i+1}Y_{\ell-i}}W_{Y_{\ell-i}Y_{\ell-i-1}}W_{Y_{\ell-j+1}Y_{\ell-j-1}}}\right)\right] \\
&= D(Q_{\ell-i}\|Q_{\ell-j}).
\end{aligned}$$

Then we note that

$$\begin{aligned}
D(Q_0\|Q_1) &= \sum_{y^\ell}\frac{1}{2}\prod_{i=2}^{\ell}W_{y_iy_{i-1}}\log\left(2W_{y_2y_1}\right) \\
&= 1-h(1/2-\epsilon),
\end{aligned}$$

where $h(\cdot)$ is the binary entropy function. Similarly

$$D(Q_1\|Q_0) = -\frac{1}{2}\log\left(1-4\epsilon^2\right).$$

For any $1 < i < \ell$,

$$\begin{aligned}
D(Q_0\|Q_i) &= \mathbb{E}_{Q_0}\left[\log\left(\frac{2W_{Y_iY_{i-1}}W_{Y_{i+1}Y_i}}{W_{Y_{i+1}Y_{i-1}}}\right)\right] \\
&= 1+\left(\frac{1}{2}+2\epsilon-\epsilon^2\right)\log\left(\frac{1}{2}+\epsilon\right) \\
&\quad +\left(\frac{1}{2}-2\epsilon+\epsilon^2\right)\log\left(\frac{1}{2}-\epsilon\right).
\end{aligned}$$

Similarly,

$$D(Q_i\|Q_0) = -\left(\frac{1}{2}-\epsilon\right)\log\left(\frac{1}{2}+\epsilon\right)+\left(\frac{1}{2}+\epsilon\right)\log\left(\frac{1}{2}-\epsilon\right).$$

Having computed these distances, we make one additional observation. For $1 \le i, j \le \ell$, $i \ne j$,

$$\begin{aligned}
D(Q_i\|Q_j) &= \mathbb{E}_{Q_i}\left[\log\left(\frac{W_{Y_{i+1}Y_{i-1}}}{W_{Y_{i+1}Y_i}W_{Y_iY_{i-1}}}\right)+\log\left(\frac{W_{Y_jY_{j-1}}W_{Y_{j+1}Y_j}}{W_{Y_{j+1}Y_{j-1}}}\right)\right] \\
&= D(Q_0\|Q_j)+D(Q_i\|Q_0).
\end{aligned}$$

Since ϵ is bounded, $D(Q_0\|Q_i) = O(1)$ and $D(Q_i\|Q_0) = O(1)$. This in turn proves that the KL divergences between any two hypotheses is a constant independent of ℓ. ∎

REFERENCES

[1] A. Kittur, E. H. Chi, and B. Suh, "Crowdsourcing user studies with Mechanical Turk," in *Proc. SIGCHI Conf. Hum. Factors Comput. Syst. (CHI 2008)*, Apr. 2008, pp. 453–456.

[2] P. G. Ipeirotis, F. Provost, and J. Wang, "Quality management on Amazon Mechanical Turk," in *Proc. ACM SIGKDD Workshop Human Comput. (HCOMP'10)*, Jul. 2010, pp. 64–67.

[3] N. Scheiber, "A middle ground between contract worker and employee," *The New York Times*, Dec. 2015.

[4] F. Gino and G. Pisano, "Toward a theory of behavioral operations," *Manuf. Service Oper. Manag.*, vol. 10, no. 4, pp. 676–691, Fall 2008.

[5] D. R. Karger, S. Oh, and D. Shah, "Budget-optimal task allocation for reliable crowdsourcing systems," *Oper. Res.*, vol. 62, no. 1, pp. 1–24, Jan.-Feb. 2014.

[6] V. Misra and T. Weissman, "Unsupervised learning and universal communication," in *Proc. 2013 IEEE Int. Symp. Inf. Theory*, Jul. 2013, pp. 261–265.

[7] H. J. Jung, Y. Park, and M. Lease, "Predicting next label quality: A time-series model of crowdwork," in *Proc. AAAI Conf. Human Comput. and Crowdsourcing (HCOMP'14)*, Nov. 2014, pp. 87–95.

[8] D. R. Karger, S. Oh, and D. Shah, "Efficient crowdsourcing for multi-class labeling," in *Proc. ACM SIGMETRICS Int. Conf. Meas. Model. Comput. Syst.*, Jun. 2013, pp. 81–92.

[9] N. B. Shah, S. Balakrishnan, and M. J. Wainwright, "A permutation-based model for crowd labeling: Optimal estimation and robustness," arXiv:1606.09632, Jun. 2016.

[10] A. Vempaty, L. R. Varshney, and P. K. Varshney, "Reliable crowdsourcing for multi-class labeling using coding theory," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 4, pp. 667–679, Aug. 2014.

[11] L. R. Varshney, P. Jyothi, and M. Hasegawa-Johnson, "Language coverage for mismatched crowdsourcing," in *Proc. 2016 Inf. Theory Appl. Workshop*, Feb. 2016.

[12] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh, "Clustering with Bregman divergences," *J. Mach. Learn. Res.*, vol. 6, pp. 1705–1749, Oct. 2005.

[13] V. Misra, "Universal communication and clustering," Ph.D. dissertation, Stanford University, Jun. 2014.

[14] C. Chan, A. Al-Bashabsheh, J. B. Ebrahimi, T. Kaced, and T. Liu, "Multivariate mutual information inspired by secret-key agreement," *Proc. IEEE*, vol. 103, no. 10, pp. 1883–1913, Oct. 2015.

[15] N. Slonim, N. Friedman, and N. Tishby, "Agglomerative multivariate information bottleneck," in *Advances in Neural Information Processing Systems 14*, T. G. Dietterich, S. Becker, and Z. Ghahramani, Eds. Cambridge, MA: MIT Press, 2002, pp. 929–936.

[16] J. Zhang, V. S. Sheng, J. Wu, and X. Wu, "Multi-class ground truth inference in crowdsourcing with clustering," *IEEE Trans. Knowl. Data Eng.*, vol. 28, no. 4, pp. 1080–1085, Apr. 2016.

[17] K. Nagano, Y. Kawahara, and S. Iwata, "Minimum average cost clustering," in *Advances in Neural Information Processing Systems 23*, J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, Eds. MIT Press, 2010, pp. 1759–1767.

[18] Q. Li, A. Vempaty, L. R. Varshney, and P. K. Varshney, "Multi-object classification via crowdsourcing with a reject option," arXiv:1602.00575 [cs.LG]., Jun. 2016.

[19] S. M. Ali and S. D. Silvey, "A general class of coefficients of divergence of one distribution from another," *J. R. Stat. Soc. Ser. B. Methodol.*, vol. 28, no. 1, pp. 131–142, 1966.

[20] I. Csiszár, "Information-type measures of difference of probability distributions and indirect observations," *Stud. Sci. Math. Hung.*, vol. 2, pp. 299–318, 1967.

[21] S. S. Dragomir, Ed., *Inequalities for Csiszár f-Divergence in Information Theory*, ser. RGMIA Monographs. Victoria University, 2000.

[22] T. Kailath, "The divergence and Bhattacharyya distance measures in signal selection," *IEEE Trans. Commun. Technol.*, vol. COM-15, no. 1, pp. 52–60, Feb. 1967.

[23] I. Csiszár and Z. Talata, "Context tree estimation for not necessarily finite memory processes, via BIC and MDL," *IEEE Trans. Inf. Theory*, vol. 52, no. 3, pp. 1007–1016, Mar. 2006.

[24] B. Yu, "Assouad, Fano, and Le Cam," in *Festschrift for Lucien Le Cam: Research Papers in Probability and Statistics*, D. Pollard, E. Torgersen, and G. L. Yang, Eds. New York: Springer, 1997, pp. 423–435.

[25] G. Valiant and P. Valiant, "Estimating the unseen: An $n/\log(n)$-sample estimator for entropy and support size, shown optimal via new CLTs," in *Proc. 43rd Annu. ACM Symp. Theory Comput. (STOC'11)*, Jun. 2011, pp. 685–694.

[26] A. Antos and I. Kontoyiannis, "Convergence properties of functional estimates for discrete distributions," *Rand. Str. & Alg.*, vol. 19, no. 3-4, pp. 163–193, 2001.

[27] L. Paninski, "Estimation of entropy and mutual information," *Neural Comput.*, vol. 15, no. 6, pp. 1191–1253, Jun. 2003.

[28] P. Netrapalli, S. Banerjee, S. Sanghavi, and S. Shakkottai, "Greedy learning of Markov network structure," in *Proc. 48th Annu. Allerton Conf. Commun. Control Comput.*, Sep. 2010, pp. 1295–1302.